arXiv:1510.07095v1 [cs.PL] 24 Oct 2015

On the Value and Limits of Multi-level Energy Consumption Static Analysis for Deeply Embedded Single and Multi-threaded Programs

Kyriakos Georgiou, Steve Kerrison, Kerstin Eder

University of Bristol

Abstract. There is growing interest in lowering the energy consumption of computation. Energy transparency is a concept that makes a program's energy consumption visible from software to hardware through the different system layers. Such transparency can enable energy optimizations at each layer and between layers, and help both programmers and operating systems make energy aware decisions. The common methodology of extracting the energy consumption of a program is through direct measurement of the target hardware. This usually involves specialized equipment and knowledge most programmers do not have. In this paper, we examine how existing methods for static resource analysis and energy modeling can be utilized to perform Energy Consumption Static Analysis (ECSA) for deeply embedded programs. To investigate this, we have developed ECSA techniques that work at the instruction set level and at a higher level, the LLVM IR, through a novel mapping technique. We apply our ECSA to a comprehensive set of mainly industrial benchmarks, including single-threaded and also multi-threaded embedded programs from two commonly used concurrency patterns, task farms and pipelines. We compare our ECSA results to hardware measurements and predictions obtained based on simulation traces. We discuss a number of application scenarios for which ECSA results can provide energy transparency and conclude with a set of new research questions for future work.

1 Introduction

A substantial amount of effort has been invested into predicting the execution time of a program. However, there is little in the complementary area of energy consumption. Such information can be of significant value during the development and life time of critical systems. For example, energy consumption information can be crucial for devices that depend on unreliable, limited sources of power such as energy harvesters. Giving consideration to the energy consumption of a system at development time can avoid potential system failures due to inadequate energy supply at runtime. For systems that operate on a battery, this can provide a good approximation of the time frame in which the battery needs replacement.

The energy consumption of a program on specific hardware can always be determined through physical measurements. Although this is potentially the most accurate method, it is often not easily accessible. Measuring energy consumption can involve sophisticated equipment and special hardware knowledge. Custom modifications may be needed to probe the power supply. These conditions make it very difficult for the majority of software developers to assess a program's energy consumption.

Static Resource Analysis (SRA) provides an alternative to measurement. Significant progress has been made in the area of Worst Case Execution Time (WCET) prediction using static techniques that determine safe upper bounds for the execution time of programs. This naturally leads to the question of whether similar techniques can be used to bound the energy consumption of programs, and, if so, how effective they can be. A popular approach used for WCET is the Implicit Path Enumeration Technique (IPET), which retrieves the worst case control flow path of programs based on a timing cost model. Instead, in [1], an energy model that assigns energy values to blocks of Instruction Set Architecture (ISA) code is used, and the authors claim to statically estimate Worst Case Energy Consumption (WCEC).

However, in contrast to timing, energy consumption is data sensitive, i.e. the energy cost of executing an instruction varies depending on (the circuit switching activity caused by) the operands used. This effect is not captured in non data sensitive energy models, i.e. models that assign a single energy consumption value to each entity, e.g. to each instruction. Such models typically are characterized based on averages obtained from measuring the energy consumed when random data is being processed [2]. Alternatively, the highest energy consumption measured could be used for model characterization. As a consequence, when a non data sensitive energy model is used, the safety of the bounds retrieved from a worst case path static analysis might be undermined by worst case data scenarios for models that provide average energy consumption costs. On the other hand, the use of worst case models is known to lead to over-estimations [3] affecting the tightness of the retrieved bounds, because it is unlikely that the data that triggers the worst case energy consumption for one instruction also does this for all subsequent instructions in a program. This problem applies to all previous works that perform static analysis for energy consumption, as they combine non data sensitive bound analysis techniques with non data sensitive energy models. In [1] static analysis for WCEC is claimed by maximising the switching activity factor for each simulated component. However, the model abstraction level used does not guarantee that a physical implementation would behave in this way. We use a model with a similar constraint, where the data input that would trigger the worst case per instruction is not known, and so cannot assert the results to be WCEC.

In this paper we thoroughly investigate the value and limitations of using IPET in combination with non data sensitive energy models to perform Energy Consumption Static Analysis (ECSA) in the context of deeply embedded hardware, in our case the XMOS XS1-L "Xcore" [4]. The Xcore is a multi-threaded

deeply embedded processor with time-deterministic instruction execution. Such systems are simpler than general purpose processors and favor predictability and low energy consumption over maximizing performance. The absence of performance enhancing complexity at the hardware level, such as caches, provides us with an ideal setting to evaluate ECSA.

We base our investigation on an ISA-level multi-threaded energy model for the Xcore [5]. This model was characterized using constrained pseudorandom input data and associates a single averaged energy cost with each instruction in the XMOS ISA. We refined this model to one that is well suited for ECSA as it represents both static and dynamic power contributions to better reflect inter-instruction and inter-thread overheads; this improved model accuracy by an average of 4%. In addition to using this model for ECSA, we also used it to compare ECSA results with predictions based on statistics obtained from simulation traces.

For our study we have developed an IPET-based ECSA, which we use together with the non data sensitive ISA-level energy model described above, to predict the energy consumption of single and multi-threaded programs. With respect to the latter we focus on two commonly used concurrency patterns in embedded programs, task farms and pipelined programs with evenly distributed workloads across threads. In addition, we have developed a novel mapping technique to lift our ISA-level energy model to a higher level, the intermediate representation of the compiler, namely LLVM IR [6], implemented within the LLVM tool chain [7]. This enables ECSA to be performed at a higher level than ISA, thus introducing energy transparency into the compiler tool chain by making energy consumption information accessible directly to the optimizer.

Performing ECSA on multi-threaded programs and at the LLVM IR allows a comprehensive analysis of the energy consumption predictions that can be obtained using this technique. Our ECSA technique is evaluated using a set of single- and multi-threaded benchmarks, mainly selected from a number of industrial embedded applications. Our results show that accurate energy estimations can be retrieved at the ISA level. The mapping technique allowed for energy consumption transparency at the LLVM IR level, with accuracy keeping within 1% of ISA-level estimations in most cases. The main contributions of this paper are:

1. Modeling the target architecture to capture its behavior statically, including refinement of an existing ISA-level energy model, improving its accuracy by around 4% (Section 3.1);
2. Formalization and implementation of a novel mapping technique that lifts an ISA-level energy model to a higher level, the intermediate representation of the LLVM compiler, which allows ECSA of programs at the LLVM IR level (Section 3.2);
3. ECSA on a set of multi-threaded programs (Section 3.5), focusing on task farms and pipelines, two commonly used concurrency patterns in embedded computing;

4. Comprehensive evaluation of our ECSA on a set of industrial benchmarks and detailed analysis of results (Section 4.2);
5. Discussion of the practical value and limitations of how such analysis can be useful for software developers, compiler engineers, development tools and Real Time Operating Systems (RTOS) (Section 4.3).

The rest of the paper is organized as follows. Section 2 critically reviews previous work on energy modeling and SRA, with a focus on SRA for energy consumption and the effects of combining non data sensitive bound analysis techniques with non data sensitive energy models. Section 3 introduces in detail the components of our analysis, in particular the formalization and implementation of our mapping technique, and how ECSA can be applied to multi-threaded programs. Our experimental evaluation methodology, benchmarks and results are presented and discussed in Section 4. Section 5 concludes the paper, outlines opportunities for future work and raises a number of research questions to stimulate further research in ECSA.

2 Background

The work presented in this paper builds upon two areas: processor energy modeling and SRA. This section establishes the background work of both.

2.1 Energy modeling of embedded processors

Energy modeling can be performed at various levels of abstraction, from gate- or transistor-level in detailed hardware simulation [8], up to high-level modeling of whole applications. Although the hardware components are responsible for power dissipation and thus consumption of energy, the behavior of that hardware is largely controlled by the software running upon it. As such, writing software that makes efficient use of the underlying hardware has been identified as the most important step in energy efficient software development [9]. For energy modeling to be useful to a software developer, models must convey information that can be related to the code the developer is writing.

The ISA is a practical level of abstraction for energy modeling of software, because it expresses the underlying operations performed by the hardware and its relationship with the intent of the software. In [2] an ISA-level energy model is proposed that obtained energy consumption data through hardware measurements of large loops of individual instructions. The total cost of a program is composed of instruction costs, inter-instruction costs (the effects of switching from one instruction to the next), and externally modeled behaviour such as activity in the memory hierarchy.

This work was initially applied to x86 and SPARC architecture processors, operating with an accuracy of within 10% of the hardware. It was extended to form a framework for architecture-level power analysis, Wattch [10]. The Sim-Panalyser [11] uses a similar approach, built on top of the SimpleScalar architecture simulation framework [12].

If additional characteristics of processor activity are considered, such as bit-flips in the data-path, a more accurate data-dependent model can be produced, such as that of [13, 14]. This requires more detailed information from simulation in order to supply additional model parameters, but has been demonstrated to bring accuracy to within 1.7% of the hardware. It is still an abstraction away from the internal switching activity of functional units, however. Observing the results in [15], some functional units may be more dependent on their internal structure than input/output Hamming weight with respect to data-dependent power.

Using similar approaches to Steinke and Tiwari, additional processor architectures such as VLIW DSPs have also been modeled, with 4.8% [16] and 1.05% [17] accuracy. Alternative approaches to modeling include representing activity in terms of the processor's functional blocks [17], energy profiling of the most commonly used software library functions [18], and construction of model parameters through linear regression [19].

In [1], a micro-architectural energy model was created, considering functional units activity, clock gating and pipeline progression for a simulated processor. This model was used for WCEC static analysis. To retrieve safe bounds, the switching activity factor was set to the maximum, 1.0, for each component. This led to significant energy consumption over-estimations in some cases, up to 33%, and assumes that the model accurately reflects a physical implementation.

In architectures where performance counters are available, these can be used to characterize the processor energy consumption based on the conditions affecting these counters, such as cache misses and pre-fetches. Simulations that model these performance events can then be used to predict the energy consumption of an application. This has been applied to processors of various levels, from embedded XScale [20] to Xeon Phi accelerators [21].

The discussed approaches achieve varying levels of accuracy, all within a 10% error margin. The comparison points vary between methods, so the accuracies are not necessarily directly comparable. However, the prior work motivates new models to achieve a similar margin. In many of the above examples, the models target a 'typical' energy characterization, where the modeled energy consumption is based on random or non-exhaustive input data sets. For a given application, some additional error margin will be introduced based on the particular characteristics of its dataset. This forms a part of the model error, in addition to the errors arising from the abstractions applied in each model type. The work presented in this paper, which examines multiple abstraction levels, seeks to identify each point at which inaccuracies may be introduced into the estimation process. This is important to assesses usefulness of estimations produced by static analysis, and will be discussed in Section 4.2

2.2 Static Resource Analysis

SRA is a methodology to determine the usage of a resource (usually time or energy or both) for a specific task when executed on a piece of hardware, without actually executing the task. This requires accurate modeling of the hardware's

behavior in order to capture the dynamic functional and non-functional properties of task execution. Determining these properties accurately is known to be undecidable in general. Therefore, to extract safe values for the resource usage of a task, a sound approximation is needed [22, 23].

SRA has been mainly driven by the timing analysis community. Static cost analysis techniques based on setting up and solving recurrence equations date back to Wegbreit's [24] seminal paper, and have been developed significantly in subsequent work [25, 26, 27, 28, 29, 30]. Other classes of approaches to cost analysis use dependent types [31], SMT solvers [32], or size change abstraction [33].

For performing an accurate WCET static analysis, there are four essential components [22]:

1. Value analysis: mainly used to analyze the behavior of the data cache.
2. Control flow analysis: used to identify the dynamic behavior of a program.
3. Low level or processor behavior analysis: attempts to retrieve timing costs for each atomic unit on a given hardware platform, such as an instruction or a basic block (BB) in a Control Flow Graph (CFG) for a processor.
4. Calculation: uses the results from the two previous components to estimate the WCET. Most common techniques used for calculation of the WCET are the IPET, the path-based techniques and the tree-based methods [34].

Three of the above components, namely the control flow analysis, low level analysis and calculation, are adopted in our work and will be further explained in Section 3.

IPET is one of the most popular methods used for WCET analysis [35, 36, 37, 34, 38]. In this approach, the CFG of a program is expressed as an Integer Linear Programming (ILP) system, where the objective function represents the execution time of the program. The problem then becomes a search for the WCET by maximizing the retrieved objective function under some constraints on the execution counts of the CFG's basic blocks. The main advantage of this technique is the ability to determine the basic blocks in the worst case execution path and their respective execution counts without the need to extract the explicit worst execution path (ordered list of the executed basic blocks). This is more efficient than path based techniques for retrieving WCET bounds [34].

In the presence of caches or a complex processor pipeline, the ILP solving complexity can increase dramatically, making IPET not practical for WCET. Abstract interpretation [39], a technique used to facilitate data flow analysis, can then be used in conjunction with IPET to allow WCET in such cases [36].

Although significant research has been conducted in static analysis for the execution time estimation of a program, there is little on energy consumption. One of the few approaches [40] seeks to statically infer the energy consumption of Java programs as functions of input data sizes, by specializing a generic resource analyzer [29, 41] to Java bytecode analysis [42]. However, a comparison of the results to actual measurements was not performed. Later, in [43], the same generic resource analyzer was instantiated to perform energy analysis of XC programs [44] at the ISA level based on ISA-level energy models and including a comparison to actual hardware measurements. However, the scope of this

particular analysis approach was limited to a small set of simple benchmarks because information required for the analysis of more complex programs, such as program structure and types, is not available at the ISA level. The analysis presented in this paper does not rely on such information. A similar approach, using cost functions, was used in [45]. The analysis was performed at the LLVM IR level, using the mapping technique that we formalise and describe in full detail for the first time in this paper. Although the range of programs that could be analyzed was improved compared to [43], the complexity of solving recurrence equations for analysing larger programs proved a limiting factor.

In [1] the WCEC for a program was inferred by using the IPET first introduced in [35]. They claim WCEC analysis, and experimental results indicate that all energy estimations over-approximate the energy consumptions retrieved from simulation. However, infeasible paths were not excluded from analysis, and there is no guarantee that the comparison test cases used in simulation were the actual worst cases.

Similarly, in [46] the authors attempt to perform static worst case energy consumption analysis for a simple embedded processor, the `Cortex M0+`. This analysis is also based on IPET combined with a so called absolute energy model, an energy model that is said to provide the "maximum energy consumption of each instruction". The authors argue that they can retrieve a safe bound. However, this is demonstrated on a single benchmark, `bubblesort`, only. The bound is 19% above a single hardware measurement; the authors acknowledge that this approach leads to over-approximations. Furthermore, the hardware measurement used as a base line to evaluate the prediction obtained from static analysis only captures the algorithm's worst case complexity scenario, no information is given on the actual data to provide insight into the effect of data switching activity on energy consumption. This can be misleading, since two sets of different input data might have the same algorithmic worst case behavior, but can be very different with respect to their total energy consumption. In practice, this gives rise to a range of energy consumption measurements for different input data all triggering the algorithmic worst case path. For instance, for the Xcore architecture the energy consumption of the `MatMult 4` threads benchmark [47] for the same size of matrices, ranges from 4.1 to 4.9 nJ depending on the used data. We have closely investigated this and discuss our findings in Figure 5.

All of the reviewed previous works for static energy consumption analysis used worst case path analysis methods combined with non input pattern-depended energy models. Currently, there is no practical method to perform average case static analysis [48]. One of the most recent works towards average case SRA, demonstrates that compositionality combined with the capacity for tracking data distributions unlocks the average case analysis, but novel language features and hardware designs are required to support these properties [49]. Furthermore, developing a data sensitive energy model requires detailed knowledge of and access to the RTL, since the power dissipation is highly depended on the switching activity inside the circuits [50]; this is a challenge in itself. This situation has motivated us to conduct a comprehensive study to fully understand the

value and limitations of ECSA, using IPET-based analysis in combination with a single cost energy model, for both single and multi-threaded code at the ISA and LLVM IR levels of abstraction.

3 Energy Consumption Static Analysis

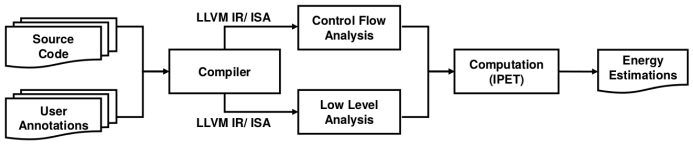

Fig. 1: Overview of our energy consumption static analysis.

Figure 1 shows the ECSA process for both, analysis at ISA and LLVM IR level. The source code together with any user annotations (e.g. to provide loop bounds) is sent to the compiler which emits the LLVM IR and the ISA code. Low level analysis, analysis of program control flow and computation of the energy consumption estimations is then applied on both levels. For the LLVM IR analysis an extra step is required at the compilation phase for the energy characterization of the LLVM IR instructions as detailed in Section 3.2. In the rest of this section we briefly introduce each ECSA stage.

3.1 Low Level Analysis

This stage aims to model the micro-architecture dynamic behavior of the processor based on an ISA-level energy model.

XMOS Xcore ISA level Energy Modeling The Xcore processor is hardware multi-threaded, providing inter-thread communication and I/O port control directly in the ISA. It is event-driven; busy waiting is avoided in favor of hardware scheduled idle periods. This makes the Xcore well suited to embedded applications requiring multiple hardware interfaces with real-time responsiveness.

The underlying energy model for this work is captured at the ISA level. Individual instructions from the ISA are assigned a single cost each. These can then be used to compute power or energy for sequences of instructions. The model also captures the cost of thread scheduling performed by the hardware, in accordance with a series of profiling tests and measurements, because it influences the energy consumption of program execution. Instructions from runnable threads are scheduled round-robin by the hardware. To avoid data hazards, the processor's four stage pipeline may only contain one instruction from each thread. If the number of runnable threads is less than four, there will be empty pipeline stages.

The modeling technique is built upon [2], as discussed in Section 2.1, which is adapted and extended to consider the scheduling behavior and pipeline characteristics of the Xcore [5]. A new version of this model that is well suited for static analysis has been developed. It represents energy in terms of static and dynamic power components to better reflect inter-instruction and inter-thread overheads. This has improved model accuracy by an average of 4%.

$$E_{\text{prg}} = (P_s + P_{di}) \cdot T_{\text{idl}} + \sum_{i \in \text{prg}} \left(\frac{P_s + P_i M_{N_p} O}{N_p} \cdot 4 \cdot T_{\text{clk}} \right), \text{ where } N_p = \min(N_t, 4) \tag{1}$$

In Equation (1), E_{prg} is the energy of a program, formed by adding the energy consumed at idle to the energy consumed by every instruction, i, executed in the program. At idle, only a base processor power, the sum of its static, P_s, and dynamic idle power, P_{di}, is dissipated for the total idle time, T_{idl}. For each instruction, static power is again considered, with additional dynamic power for each particular instruction, P_i. The dynamic power contribution is then multiplied by a constant inter-instruction overhead, O, that has been established as the average overhead of instruction interleaving. This is then multiplied by a scaling factor to account for the number of threads in the pipeline, M_{N_p}. The result is divided by the number of instructions in the pipeline, which is at most four and is dependent upon the number of active threads, N_t. Each instruction completes in four cycles, so $4 \cdot T_{\text{clk}}$ gives the energy contribution of the given instruction, based on the calculated power.

When more than four threads are active, the issue rate of instructions per thread will be reduced. The energy model accounts for this with the min term in Equation (1). From a purely timing perspective, the latency between instruction issues for a thread is $\max(N_t, 4) \cdot T_{\text{clk}}$. This property means that instructions are time-deterministic, provided the number of active threads is known. A thread may stall in order to fetch the next instruction. This is also deterministic and can be statically identified [51, pp. 8–10]. These instruction timing rules have been used in simulation based energy estimation, and are also utilized in the multi-threaded static analysis performed in this paper. Both simulation and static analysis must be able to determine N_t, the number of active threads, in order to accurately estimate energy consumption.

A limited number of instructions can be exceptions to these timing rules. The divide and remainder instructions are bit-serial and take up to 32 cycles to complete. Resource instructions may block if a condition of their execution is not met, e.g. waiting on inbound communication causes the instruction's thread to be de-scheduled until the condition becomes satisfied. This paper focuses its contributions on fully predictable instructions, with timing disturbances from communication forming future work.

The cost associated with an instruction represents the average energy consumption obtained from measuring the energy consumed during instruction execution based on constrained pseudo-randomly generated operands using the setup described in [5]. Thus, this model does not explicitly consider the range of

input data values and how this may affect consumption. Empirical evidence indicates that such factors can contribute to the dynamic energy consumption [3]. The implications of using a random data constructed single value energy model with a bound SRA for ECSA are discussed in Section 4.2.

Utilizing the Xcore energy model in static analysis To determine the energy consumption of a program based on Equation (1) the program's instruction sequence, $\langle i_1, \ldots, i_n \rangle$, the idle time T_{idl}, and the number of active threads N_p during instruction execution must be known. In [5] Instruction Set Simulation (ISS) was used to gather full trace data or execution statistics to obtain these parameters. In this work we use ISS only as a reference for comparison of ECSA results, with a second reference being direct hardware measurement.

ECSA thus needs to extract the CFGs for each thread and identify the interleavings between them. This allows for each instruction in the program to identify the N_p component in Equation (1). It also allows to estimate the total idle time, T_{idl} of the program. For single-threaded programs the energy characterization of the CFG is straightforward as there is no thread interleaving. The IPET can be directly applied on the energy characterized CFG to extract a path that bounds the energy consumption of the program, as described in Section 3.4. For arbitrary multi-threaded programs, energy characterizing the CFGs of each thread using static analysis is challenging. We have therefore concentrated on two commonly used concurrency patterns, task farms and pipelines, which we use with evenly distributed workloads across threads.

In addition to the instructions defined in the ISA, a Fetch No-Op (FNOP) can also be issued by the processor. These occur deterministically [51, pp. 8–10]. FNOPs can have a significant impact on energy consumption, particularly within loops. To account for FNOPs in static analysis, the program's CFG at ISA level is analyzed. An instruction buffer model is used to determine where FNOPs will occur in a basic block. However, one particular FNOP case is dependent on the dynamic branching behavior of the program, in which case we overestimate FNOPs. Further implementation details on FNOPs modeling can be found in [52].

3.2 Mapping an ISA Energy Model to LLVM IR

Although substantial effort has been devoted to ISA energy modeling, there is little research into modeling at higher levels of program representation, where precision can decrease. In [53], statistical analysis and characterization of LLVM IR code is performed. This is combined with instrumentation and execution on a target host machine to estimate the performance and energy requirements in embedded software. Transferring the LLVM IR energy model to a new platform requires performing the statistical analysis again. The mapping technique we present here is fully portable. It requires only the adjustment of the LLVM mapping pass to the new architecture. Furthermore, our LLVM IR mapping technique provides on-the-fly energy characterization that allows to take into

consideration the compiler behavior, CFG structure, types and other aspects of instructions.

Formal specification of the mapping Our mapping technique determines the energy characteristics of LLVM IR instructions. Thus, mapping links LLVM IR instructions with machine specific ISA instructions. ISA level energy models can then be propagated up to LLVM IR level, allowing energy consumption estimation of programs at that level. We formalize the mapping as follows. For a program P, let

$$\text{IRprog}_L = \{1, 2, ..., n\} \tag{2}$$

be the ID numbers of P's LLVM IR instructions and therefore

$$\text{IRprog} = \langle ir_1, ir_2, ..., ir_n \rangle \tag{3}$$

is the sequence of LLVM IR instructions for P.

$$\text{T}_{arch}(\text{IRprog}) = \text{ISAprog} \tag{4}$$

is an architecture specific compiler back end that can translate the IRprog to

$$\text{ISAprog} = \langle (isa_1, m_1), (isa_2, m_2), ..., (isa_k, m_l) \rangle \text{ where } m_1, m_2, ..., m_l \in \text{IRprog}_L \tag{5}$$

which is the sequence of ISA instructions for P, together with the ID of the LLVMIR instruction from which each isa_k originated. If an isa_k comes from more than one LLVM IR instructions, then T_{arch} chooses the ID of one of them to assign to isa_k.

$$\text{M}(ir_i) = \{isa_j | ir_i \in \text{IRprog} \wedge \text{ISAprog} = \text{T}_{arch}(\text{IRprog}) \wedge (isa_j, i) \in \text{ISAprog}\} \text{ and } \forall\ ir_n, ir_k \in \text{IRprog} \wedge ir_n \neq ir_k \text{ then } \text{M}(ir_n) \cap \text{M}(ir_k) = \emptyset \tag{6}$$

is a mapping function that captures a 1:m relation from IRprog to ISAprog instructions. Therefore,

$$\text{E}(ir_i) = \sum_{isa_j \in S} \text{E}(isa_j) \text{ where } ir_i \in \text{IRprog} \wedge isa_j \in \text{ISAprog} \wedge \text{S} = \text{M}(ir_i) \tag{7}$$

represents the energy consumption of an LLVM IR instruction as the sum of the energy consumed by all ISA instructions associated with that LLVM IR instruction.

By instantiating the above mapping to a specific architecture, LLVM IR energy characterization can be retrieved. The accuracy of this characterization can vary for different architectures. If the accuracy is not adequate, then a tunning phase can be introduced to account for any specific compiler or architecture behavior. An example of such tunning is given in the next section, which accounts for *phi-nodes* and FNOPs.

Xcore mapping instantiation and tuning In our case, the T_{arch} function is the XMOS tool chain lowering phase that translates the LLVM IR to Xcore specific ISA. Our mapping implementation leverages the debug mechanism in the XMOS compiler tool chain, in order to enable T_{arch} to assign to each ISA instruction the ID of the LLVM IR instruction it originated from. This is typically used by the programmer to identify and fix problems in application code. Debug symbols are created during compilation to assist with this. These symbols are propagated to all intermediate code layers and down to the ISA code. Debug symbols can express which programming language constructs generated a specific piece of machine code in a given executable module. In our case, these symbols are generated by the front end of the XMOS compiler in standard DWARF format [54]. These are transformed to LLVM metadata [55] and attached to the LLVM IR.

During the lowering phase of compilation, LLVM IR code is transformed to a target ISA by the back end of the compiler, with debug information stored alongside it as LLVM metadata. Naturally, the accuracy of debug information in the output executable is reduced if the number of optimization passes is increased. This is due to portions of the initial LLVM IR either being discarded or merged during these passes.

Tracking this information gives an $n : m$ relationship between instructions at the different layers, because source code instructions can be translated to many LLVM IR instructions, and these again into many ISA instructions. This $n : m$ relation prevents ECSA from providing accurate energy values and therefore the mapping introduced in Section 3.2, requires Equation (6) to create an 1:m relation between the LLVM IR and ISA code.

To address this issue, we created an LLVM pass that traverses the LLVM IR and replaces source location information with LLVM IR location information. The location information represents the IRprog_L in Equation (2). The LLVM pass runs after all optimization passes, just before emitting ISA code. The optimized LLVM IR is closer in structure to the ISA code than the unoptimized version. Using this method a $1 : m$ mapping between LLVM IR instructions and ISA instructions can be extracted by Equation (6). Once the mapping has been performed for a program, the energy values for groups of ISA instructions are aggregated and then associated with their single corresponding LLVM IR instruction using Equation (7).

An example mapping is given in Figure 2. On the left hand side is a part of the LLVM IR CFG of a program, which represents the IRprog in Equation (3), along with the debug location, LLVMIR_L in Equation (2), for each LLVM IR instruction. The right hand side shows the corresponding ISA CFG, together with the debug locations for each ISA instruction, given by T_{arch}. The coloring of the instructions demonstrates the mapping between the two CFGs' instructions using Equation (6). Now, one LLVM IR instruction is matched to many ISA instructions, but each ISA instruction is mapped to only one LLVM IR instruction. Some LLVM IR instructions are not mapped, because they are removed during the lowering phase of the compiler. This mapping also guarantees that all

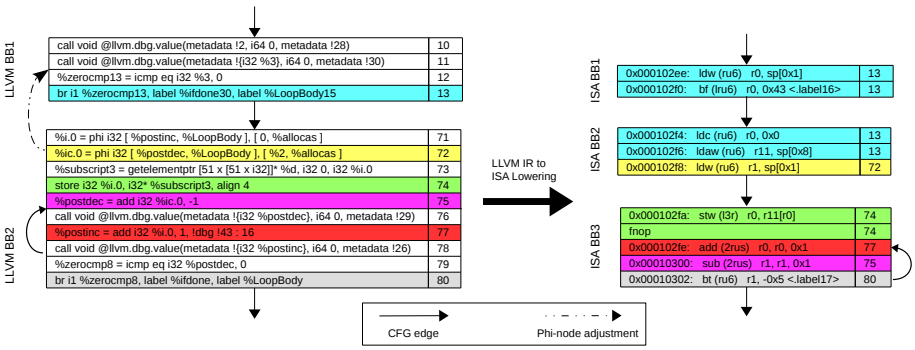
LLVM BB1
call void @llvm.dbg.value(metadata !2, i64 0, metadata !28) 10
call void @llvm.dbg.value(metadata !{i32 %3}, i64 0, metadata !30) 11
%zerocmp13 = icmp eq i32 %3, 0 12
br i1 %zerocmp13, label %ifdone30, label %LoopBody15 13
LLVM BB2
%i.0 = phi i32 [%postinc, %LoopBody], [0, %allocas] 71
%ic.0 = phi i32 [%postdec, %LoopBody], [%2, %allocas] 72
%subscript3 = getelementptr [51 x [51 x i32]]* %d, i32 0, i32 %i.0 73
store i32 %i.0, i32* %subscript3, align 4 74
%postdec = add i32 %ic.0, -1 75
call void @llvm.dbg.value(metadata !{i32 %postdec}, i64 0, metadata !29) 76
%postinc = add i32 %i.0, 1, !dbg !43 : 16 77
call void @llvm.dbg.value(metadata !{i32 %postinc}, i64 0, metadata !26) 78
%zerocmp8 = icmp eq i32 %postdec, 0 79
br i1 %zerocmp8, label %ifdone, label %LoopBody 80
LLVM IR to
ISA Lowering
ISA BB1
0x000102ee: ldw (ru6) r0, sp[0x1] 13
0x000102f0: bf (lru6) r0, 0x43 <.label16> 13
ISA BB2
0x000102f4: ldc (ru6) r0, 0x0 13
0x000102f6: ldaw (ru6) r11, sp[0x8] 13
0x000102f8: ldw (ru6) r1, sp[0x1] 72
ISA BB3
0x000102fa: stw (l3r) r0, r11[r0] 74
fnop 74
0x000102fe: add (2rus) r0, r0, 0x1 77
0x00010300: sub (2rus) r1, r1, 0x1 75
0x00010302: bt (ru6) r1, -0x5 <.label17> 80
CFG edge
Phi-node adjustment

Fig. 2: Fine grained 1:m mapping including our LLVM mapping pass.

ISA instructions are mapped to LLVM IR, so there is no loss of recorded energy between the two levels.

Additional optimizations are performed during the lowering phase from LLVM IR to ISA, such as peephole optimizations and target specific optimizations. These can affect the mapping, but not to the same degree as the LLVM optimizations. A tuning phase can be introduced after the mapping, to account for them.

The mapping instantiation for the Xcore architecture was able to provide an average energy estimation deviation of 6% from the predictions on the ISA level. An additional tuning phase is introduced after the mapping, to account for specific compiler and architecture behavior. This improved the mapping accuracy, narrowing the gap between ISA and LLVM IR energy predictions to an average of 1% as discussed in 4.3.

LLVM IR *phi-nodes* are an example of such tuning. *Phi-nodes* can be introduced at the start of a BB as a side effect of the Single Static Assignment (SSA) used for variables in the LLVM IR. A *phi-node* takes a list of pairs, where each pair contains a reference to the predecessor block together with the variable that is propagated from there to the current block. The number of pairs is equal to the number of predecessor blocks to the current block. A *phi-node* can create inaccuracies in the mapping when LLVM IR is lowered to ISA code that no longer supports SSA, because it can be hoisted out from its current block to the corresponding predecessor block. For blocks in loops this can lead to a significant analysis error.

Whenever the tuning phase is able to track these cases, it can adjust the energy figures for each LLVM IR BB accordingly. An example of this is given in Figure 2 at debug location number 72. Its corresponding ISA instruction is hoisted out from the loop BB `ISA BB3` and into `ISA BB2`. This is tracked by the mapping, and the equivalent hoisting is done at LLVM IR level, thus correctly assigning energy values to each LLVM IR block. Similar errors can be introduced by branching LLVM IR instructions with multiple targets, since in the Xcore

ISA only single target branches are supported. This is also handled during the mapping phase.

As discussed in Equation (1), FNOPs can be issued by the processor and this can be statically determined at the ISA level. LLVM IR has no way to represent this. Ignoring them can therefore lead to a significant underestimation of energy at LLVM IR level. To address this, FNOPs in the lowered ISA code are assigned the debug location of an adjacent ISA instruction in the same BB by the tuning phase, thus they are accounted for in the mapped LLVM IR block.

LLVM IR instructions can be combined into a single ISA instruction. An example of such instructions are the add and multiplication ones which can be translated to the Xcore `macc` (multiply-accumulate) ISA instruction. The T_{arch} will assign the energy cost to only one of the LLVM IR instructions. Although, this is adequate for the energy characterization of LLVM IR basic blocks, if needed the tuning phase allows to associate the cost with both instructions.

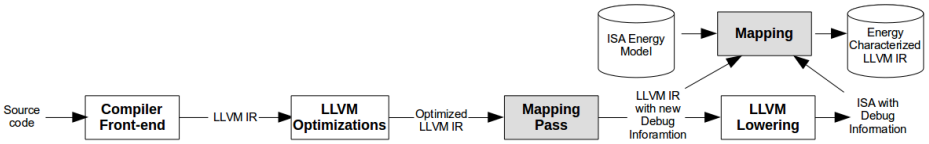

Fig. 3: Overview of the mapping process.

An overview of the mapping technique is given in Figure 3. Our mapping pass is introduced into the compilation process after LLVM optimizations. The pass also includes tuning. The mapping phase implements Relations 6 and 7. It runs after the LLVM lowering phase and maps LLVM IR instructions with the new debug locations to the emitted ISA instructions. The ISA energy model is then used to accumulate the energy value of each LLVM IR instruction based on its mapped ISA instructions.

3.3 Control Flow Analysis

This component aims to capture the dynamic behavior of the program and associates CFG BBs with the information needed for the computation step of analysis. IPET requires the CFG and call graph of a program to be constructed at the same level as the analysis. At LLVM IR level, the compiler can generate them. At ISA level a tool was created to construct them. To detect BBs that belong to a loop or recursion, we adopted and extended the algorithm in [56]. The CFGs are annotated according to the needs of the IPET described in Section 3.4. Finally, the annotated CFGs are used in the computation step to produce ILP formulations and constraints.

3.4 Computation

The IPET adopted in our work to estimate the energy consumption of a program is based on [57]. To construct the ILP system needed for IPET, we use information produced from the previous two components. The method of ILP formulation along with the constraints needed to bound the problem and optimize it's solution can be found in the seminal paper [57]. To infer the energy consumption, instead of using the time cost of a CFG basic block we are using its energy cost, as provided by the respective energy model.

Constraints are used to capture information that can affect a program's dynamic behavior, such as bounded loop iterations, or path information, such as infeasible paths. Usually, this information can only be specified by the programmer, as it depends on the program semantics and cannot be extracted by the static analysis. The minimum required user input to enable bounding of the problem is the declaration of loop bounds. This is also standard practice in timing analysis [22]. Providing this kind of information is usually easy, as the loop bounds are typically known by programmers of timing critical embedded programs. Further constraints, such as denoting infeasible paths in a CFG, can be provided to extract more accurate estimations. The user provides this information as source code annotations. The annotation language used in this work can be found in [58].

3.5 Analysis of multi-threaded programs

In this paper we present the first steps towards ECSA of multi-threaded programs. Two concurrency patterns are considered: replicated threads with no inter-thread communication, working on different sets of data (task farms), and pipelines of communicating threads. For both cases, we consider evenly distributed, balanced work loads. In the former case, an example use is simultaneously processing multiple independent data. In the latter case, pipelining enables parallelism to be used to improve performance when processing a single data stream.

There is a fundamental difference when statically predicting the case of interest (worst, best, average case) for time and for energy for multi-threaded programs. Generally, for time only the computations that contribute to the path forming the case of interest must be considered. For energy, all computations taking place during the case of interest must be considered. For instance, in an unbalanced task farm, the WCET will be equivalent to the longest running thread. To bound energy, the energy consumption of each thread needs to be aggregated. This is harder since the static analysis needs to determine the number of active threads at each point in time in order to apply the energy model from Equation (1) and characterize the CFG of each thread. Then, IPET can be applied to each thread's CFG, extracting energy consumption bounds. Aggregating these together will give a loose upper bound on the program's energy consumption, meaning that the safety of the bound cannot be guaranteed.

In our balanced task farm examples, all task threads are active in parallel for the duration of the test. Thus, the number of active threads is constant, giving a constant N_t, used to determine the pipeline occupancy scaling factor, M, in Equation (1). For balanced pipelined programs, we consider the continuous, streaming data use case, so the same constant thread count property holds. In both cases IPET can be performed on each thread's CFG and the results aggregated to retrieve the total energy consumption. In this work, core-local communication is considered, which uses the same instructions as off-core communication, but no external link energy needs to be accounted for. Therefore the core energy model provides sufficient data.

For multi-threaded programs with synchronous communications, to retrieve a WCET, IPET can be applied on a global graph, connecting the CFGs of all threads along communication edges. The communication edges can be treated by the IPET as normal CFG edges and WCET can be extracted by solving the formulated problem [38]. This will return a single worst case path across the global graph. Bounding energy in this way is not possible, as parallel thread activity over time needs to be considered. Here the task is even harder in comparison to programs without communication, as activity can be blocked if the threads' workloads are unbalanced. In this case, statically determining the number of active threads at each point in time is a hard challenge.

Although the concurrency patterns addressed here can be considered as easy targets for the ECSA, they are typical embedded use cases, and as is explained in Table 1, ECSA can provide sufficiently accurate information to enable energy aware decision making. Building on this, more complex programs will be analyzed in future work, such as unique non-communicating threads rather than replicated threads, unbalanced farm and pipeline workloads and other concurrency patterns. Such programs will feature varying numbers of active threads over the course of execution. In these cases the ECSA must be extended to perform analysis that extracts all the possible combinations of thread interleaving.

This work focuses on multi-threaded communication on a single core. However, for communicating threads, the channel communication paradigms that are used by the programs at the source code level and within the ISA can also be used in a multi-core environment, creating scope for the analysis of larger systems.

4 Experimental Evaluation

To evaluate our ECSA, a series of mainly industrial benchmarks were selected with representative test cases. Both our ECSA results and estimations from ISS using the same energy model are compared to hardware measurements. The benchmarks, evaluation methodology, results and further observations are discussed in this section.

4.1 Benchmarks

Our objective is to demonstrate the value of our ECSA for common industrial, deeply embedded applications. A complete list of all the 21 benchmarks' code and summary of their attributes, can be found in [47]. Benchmarks were compiled with `xcc` version 12 [59] at optimization level `O2`; the default for most compilers.

Deeply embedded processors do not typically have hardware support for division or floating point operations, using software libraries instead. Software implementations are usually far less efficient than their hardware equivalent, both in terms of execution time and energy consumption. The effect of these software implementations on energy consumption should be known by developers, therefore we include soft division and soft float benchmarks.

A radix-4 software divider, `Radix4Div` [60], is used. A less efficient version, `B.Radix4Div`, is added for comparison. This version omits an early return when the dividend is greater than 255. A consequence of excluding this optimization is that CFG paths become more balanced, with less variation between the possible execution paths. The effect of this on the energy consumption is discussed later in this section. For software floating point, single precision `SFloatAdd32bit` and `SFloatSub32bit` operations from [61] are analyzed.

To represent common signal processing tasks, `FIR` and `Biquad` benchmarks written for the Xcore processor [62], are analyzed. In addition, a series of open source benchmarks of core algorithmic functions were selected from the `MDH WCET` benchmark suite [63]. They were modified to work with our test harness and, in some cases, to make them more parametric to function input arguments. Some were extended to be multi-threaded task farms, where the same code runs on two or four threads. To extend our analysis to multi-threaded communication programs, we analyze pipelined versions of `FIR` and `Biquad`, each formed of seven threads. These programs are the preferred form for Xcore, as spreading the computation across threads allows the voltage and frequency of the core to be lowered, significantly reducing energy consumption with the same performance as the single threaded version.

4.2 Results Analysis

The experimental results show several features, influenced by the level of multi-threading, the properties of the benchmarks, and the levels at which ECSA and modeling are performed. In this section we examine all of these in order to determine what influences ECSA accuracy at each level, highlighting both strengths and limitations.

Figure 4a presents the error margin of using our energy model with three energy estimation techniques compared to hardware energy measurements for our benchmarks. *Trace Sim* produces instruction traces from ISS, *ISA ECSA* uses the model for static analysis at the ISA level and *LLVM IR ECSA* uses our mapping technique to apply the model and analysis at LLVM IR level. For all benchmarks with multiple test parameters, the geometric mean of the errors

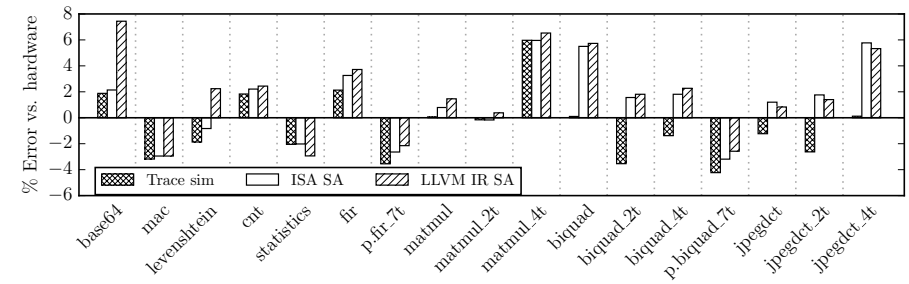

(a) All benchmarks.

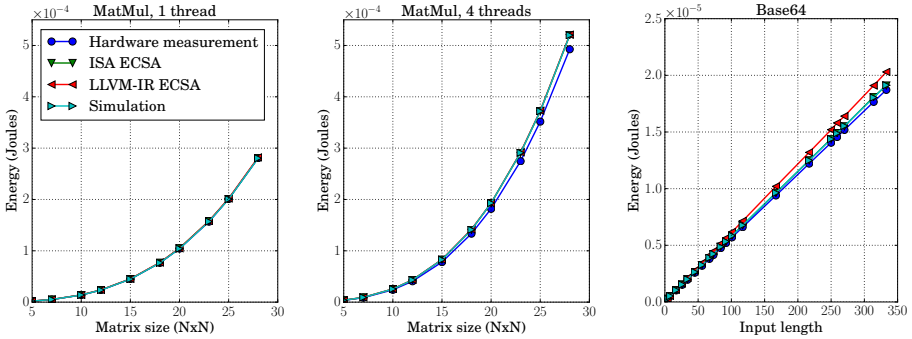

(b) Parametric benchmarks.

Fig. 4: Hardware measurements compared to ECSA and ISA trace estimation.

is used. Figure 4b compares energy estimates to hardware measurements for a range of parameters in three parametric benchmarks.

For `Levenshtein`, `MatMult 1,2,4`, `Mac`, `Cnt` and `Base64` parametric energy consumption estimations can be determined, as discussed in Section 4.3. These are expressed in terms of a function over the number of loop iterations.

The parametric benchmarks are also more data sensitive, due to the use of matrices. The hardware energy measurements for all the benchmarks using matrices were obtained by using random data to initialize them. In order to investigate the effect of different random data, the measurements were repeated 500 times for each benchmark using a different seed each time for data generation. The maximum variation observed was in the range of the measurement error, less than 0.5%, and therefore the average of these measurements was used to compare against the predicted results. The effect of using non random data will be investigated in Figure 5. For the more industry oriented benchmarks (all the `FIR`, `Biquad` and `Jpegdct` versions) real sample data where used for the hardware measurements.

For the software division and floating point benchmarks, ECSA provides a constant energy consumption upper bound across all test cases, as they con-

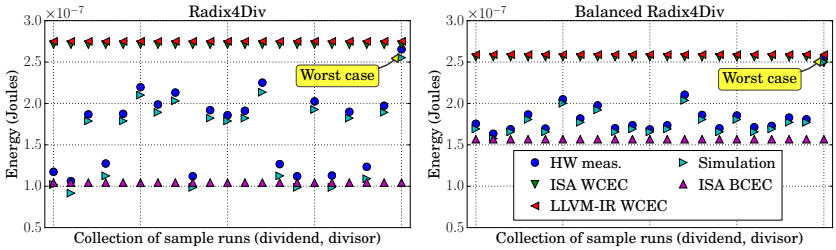

Fig. 5: Results for benchmarks with constant ECSA estimations across all test cases.

tain no loops that are directly affected by the functions' arguments. Figure 5 demonstrates this for `Radix4Div` and `B.Radix4Div`. Considering that IPET is intended to provide bounds based on a given cost model, in our case it tries to select the worst case execution paths in terms of the energy consumption. Therefore, the ECSA estimations seen in Figure 5 represent a loose upper bound on the benchmarks' energy consumption. Similar figures were also retrieved for the two `SoftFloat` benchmarks. These bounds, in most cases cannot be considered safe, as they might be undermined by the use of a non data sensitive energy model and analysis. However, they can still give the application programmer valuable guidance towards energy aware software development, as discussed in Section 4.3.

For the benchmarks in Figure 5, we sought test cases that exercise the average, best- and worst-case scenarios of each benchmark's algorithm, to compare the resultant range of energy consumption with our ECSA predictions. A good understanding of the underlying algorithms and information collected from the ISS traces was necessary to identify tests covering each scenario with certainty. This poses a challenge in guaranteeing that the cases of interest, such as worst case, have been exercised. For example, the `Radix4Div` benchmark takes two 16-bit parameters, forming a search-space of 2^{32} test cases. This was suitably small to perform an exhaustive search in order to capture the worst case empirically. However, the time cost of an exhaustive search precludes doing the same for many other benchmarks. For both `Radix4Div` variants, the upper bounds inferred by the IPET analysis are not only very close to the worst case retrieved by exhaustively searching the possible test cases, but are also safe.

Generally, for all results shown, a proportion of error is present in both forms of static analysis as well as simulation based energy estimation. The error in the ISS based estimation is a baseline for the best achievable error in static analysis, as ISS produces more accurate execution information. For all the benchmarks, the ISA ECSA results are over-approximating the trace based energy estimations. This applies also to the LLVM IR ECSA results with exception of the `statistics` benchmark. This over-approximation is a product of the bound analysis used which is trying to select the most energy costly CFG path based

on the provided cost model. A smaller difference between the ECSA results and the trace based energy estimations indicates that the execution path selected by the IPET fits better the actual execution path of a benchmark.

Measurement error analysis To assess the accuracy of ECSA predictions, reliable hardware measurements are required. We use a shunt resistor current sense circuit and data sampling hardware to obtain power dissipation with sub-milliwatt accuracy. The data capture process is explained in more detail in [5].

Measurements are subject to errors introduced through environmental factors. In particular, temperature and electrical noise can result in variations of the measured energy consumption for multiple runs of the same test. To measure the effect of these factors on our platform, we executed the `MatMult 4` thread benchmark 100,000 times. This benchmark was selected because it is particularly power intensive and likely to affect the device temperature the most. The variation observed on our hardware was less than 0.7% which we consider negligible and close to the error margin of our measurement equipment. These factors could have a more significant impact on other platforms. It is therefore important to examine them when performing ECSA.

The test harness introduces a small error by repeatedly calling the benchmark function within a loop. This is necessary to ensure an adequate number of power samples are taken during the test. However, the loop surrounding the call to the benchmark, together with the function call itself, introduce an overhead. This overhead can be significant, especially when the amount of computation in the loop body is low. To mitigate this overhead, we ensure that the loop is as efficient as possible and each benchmark sufficiently large in size. Finally, measurements were taken several times to ensure that results obtained were consistent, with less than 0.5% variation.

LLVM IR analysis accuracy This form of analysis is solely dependent on the accuracy of the mapping techniques presented in Section 3.2. As shown in Figures 4a and 5, for all benchmarks the LLVM IR ECSA results are within one percentage point error of ISA ECSA results, except for the `Base64` benchmark with a further 5.3 percentage points error. In this case the CFGs of the two levels were significantly different due to BBs introduced from branches in the ISA level CFG. This is one of the few cases where the mapper was unable to accurately track the differences between the two CFGs.

Multi-threading accuracy Three benchmarks, `MatMul`, `Biquad` and `JpegDCT`, were extended to multi-threaded versions, where each thread executes the same program and processes its own data stream. The computation performed and the energy consumption increases with the number of active threads, with a negligible change in execution time. The underlying energy model is parametric to the occupancy of the execution pipeline, which is determined by the number of running threads. As such, the estimations from the model and their relative

errors can differ when the number of threads is changed. For any given number of threads, the accuracy of the ECSA is influenced by the accuracy of the ISA level energy model.

In the case of pipelined benchmarks, `p.fir_7t` and `p.biquad_7t`, the energy model underestimates energy consumption by approximately 5%. This error is inherited by the ECSA. Further calibration of the model is required to achieve better accuracy for multi-threaded programs with communications.

Data effect Since a non data sensitive ECSA will provide a single energy estimation regardless of input data values, comparing this to hardware measurements may give a different error for each input data. To examine this, we used one of our most data sensitive benchmarks, `MatMul 4`. The smallest over-estimation when compared with the hardware measurements was 5.96% for both the ISA ECSA and trace based energy estimations. This was obtained for matrices that were initialized with randomly generated data. Since our energy model was characterized using pseudo-randomly generated data, it provides a good fit to the data used for measurement, thus this result meets our expectations. The maximum over-estimation found was 25%, by initiating both matrices with zero data, minimizing the processor's switching activity. By using the same random data in the two matrices, the over-estimation was between the two previous cases, approximately 15%. This is because the processor switching activity is less than in the case of different random data initialized matrices, and more than the case with the zero initialized matrices. Thus, users must be cautious when using ECSA with data sensitive benchmarks, as we will discuss in Section 4.3.

These findings lead to two new research questions. Firstly, for convenience, many energy models are constructed from random input data. However, as we demonstrated, the closer the data used to characterize the energy model fits the data of the use case, the more accurate the ECSA estimation. For example, `MatMul` and `fdct` are heavily used in video processing applications with highly correlated data between frames in the video stream. Therefore, a random data constructed energy model for these applications may not be suitable. How can we construct energy models that are more fit for purpose? Secondly, if a data sensitive energy model were to be constructed, how would this model be composed to be useful for ECSA? These two research questions motivate future work in this area.

Static analysis limitations ECSA suffers from all the static analysis limitations that the timing analysis faces [64]. Many of the techniques used by the timing analysis community to tackle these limitations can also be adopted in ECSA. For example, infeasible paths can lead to unrealistic estimations in both cases, energy and time. Techniques such as symbolic execution [46] used in timing analysis to exclude infeasible paths, can be also used for ECSA. For this paper, source code annotations were translated to ILP constraints, in order to exclude infeasible paths from ECSA.

As already identified, ECSA can be more complicated than timing analysis. In Figure 5, we discussed that energy consumption is sensitive to the data related switching activity in the processor, which time is not affected from. In Section 3.5 we discussed, that for multi-threaded programs, timing analysis is considered only with a single path across all the threads, but ECSA has to consider all computations active during the case of interest.

In summary, the results show that static analysis, both at ISA and LLVM-IR level, can deliver practical energy consumption estimates for a good range of single and multi-threaded programs. The estimation error for both static and simulation based techniques can be reduced if the accuracy of the underlying energy model is improved.

4.3 ECSA applications

Precise energy measurements are often not easily accessible, requiring extra equipment and hardware knowledge as well as modifications to the target hardware. This makes it very difficult for most programmers to assess a program's energy consumption. ECSA overcomes these obstacles by providing energy transparency to users and systems with a useful level of accuracy.

Trace based energy estimation allows for a very precise estimation of energy consumption for a particular program run. The program is executed in simulation with a given set of input parameters. The exact sequence of instructions can be recorded during simulation and then used to estimate energy consumption. However, a change to the input may produce a new execution path, requiring a new simulation run to extract the correct instruction sequence. Simulation is typically several orders of magnitude slower than hardware execution, making repeated simulations undesirable as a means for tuning or optimizing a program. ECSA does not depend on repeated simulation. It does not require trace data in order to provide an energy estimation. This allows for much faster estimation of a program's energy consumption.

The main difference between energy measurements, trace simulation based energy predictions and ECSA, is that the first two methods estimate the cost of the actual executed path. ECSA, however, gives an upper bound based on the cost model used. Both ECSA and trace estimations rely on the accuracy of the energy model. Further, they cannot accurately account for energy due to data-sensitive switching activity. In the rest of this section we will provide a set of guidelines on how the ECSA results should be interpreted, and how they can influence energy aware decisions that can be made by software developers, compiler engineers, development tools and RTOS.

LLVM IR level ECSA The LLVM optimizer and code emitter are the natural place for compiler optimizations. Our LLVM IR analysis results demonstrate a high accuracy with a deviation in the range of 1% from the ISA ECSA. Some LLVM IR estimations may not always be as accurate as at ISA level, but they are still of value to developers. Transparency of energy consumption

at this level enables programmers to investigate how optimizations affect their program's energy consumption [65], or even help introduce new low energy optimizations [66, 67]. This is more applicable at the LLVM IR than at the ISA level, because more program information exists at that level, such as types and loop structures. Our mapping techniques and analysis framework at the LLVM IR level are applicable to any compiler that uses the LLVM common optimizer, provided that an energy model for the target architecture is available.

For some programs, indirect jumps that are introduced at the ISA level can make it impossible to extract a CFG. While this prevents ISA level ECSA, the analysis can still be performed for these programs at LLVM IR, allowing programmers to gain energy consumption insight even when ISA level analysis is not feasible.

ECSA bound use cases Given that we are using bound analysis with an energy model characterized with random input data, we must consider the ECSA estimations as loose upper bounds of the WCEC. Although, these bounds are not safe, in most cases they can provide useful information to the programmer, e.g. to determine whether or not an application is likely to exceed an available energy budget.

The modified `B.Radix4Div` benchmark avoids an early return when the dividend is greater than 255. Omitting this optimization is less efficient, but balances the CFG paths. The effect of this modification can be seen in Figure 5. The ISA level energy consumption lower bounds (the best case retrieved by IPET) are shown. In the optimized version, the energy consumption across different test cases varies significantly, creating a large range between the upper and lower energy consumption bounds. Conversely, the unoptimized version shows a lower variation, thus narrowing the margin between the upper and lower bounds, but has a higher average energy consumption.

Knowledge of such energy consumption behavior can be of value for applications like cryptography, where the power profile of systems can be monitored to reveal sensitive information in side channel attacks [68]. In these situations, ECSA analysis can help code developers to design code with low energy consumption variation, so that any potential leak of information that could be obtained from power monitoring can be obfuscated.

Parametric resource usage equations Regression analysis was applied to the ISA level static analysis results of the benchmarks `MatMult 1,2,4`, `Mac`, `Cnt` and `Base64`. The resultant upper bound equations are shown in Table 1. The second column shows the retrieved equations which return the energy consumption predictions in nano-Joules (nJ) as a function over x, as defined in the third column. `Levenshtein` is a multi-parametric energy consumption benchmark. However, the regression analysis was unable to determine a good parametric equation for it.

Parametric resource usage equations can be valuable for a programmer or user to predict energy consumption with specific parameter values. Moreover,

embedding such equations into an operating system can enable energy aware decisions for either scheduling tasks, or checking if the remaining energy budget is adequate to complete a task. If the application permits, the operating system may also downgrade the quality of service to complete the task within a lower energy budget.

Benchmark	Regression Analysis (nJ)	x
Base64	$f(x) = 19x + 94.2$	string length
Mac	$f(x) = 15x + 21.1$	length of two vectors
Cnt	$f(x) = 19.9x^2 + 5.7x + 34.6$	matrix size
MatMul	$f(x) = 12.2x^3 + 17.5x^2 + 4.7x + 33$	size of square matrices
MatMul_2T	$f(x) = 19.3x^3 + 21.4x^2 + 5.9x + 96.8$	size of square matrices
MatMul_4T	$f(x) = 22.7x^3 + 25x^2 + 6.5x + 157.7$	size of square matrices

Table 1: Benchmarks with parametric energy consumption.

Multi-threaded ECSA The first class of parallel programs to which ECSA was applied is replicated non-communicating threads. The user can make energy aware decisions on the number of threads to use, with respect to time and energy estimations retrieved by our analysis. For example, take four independent matrix multiplications on four pairs of equally sized matrices (28×28). Our analysis will show that a single thread will have an execution time of 4x the time needed to execute one matrix multiplication. However, two threads will half the execution time and decrease the energy by 54%. Four threads which will half the execution time again, and decrease the energy by 41% compared to the two-thread version. Using more threads increases the power dissipation, but the reduction in execution time saves energy on the platform under investigation. Although there is a different estimation error between different numbers of active threads, the error range of 6% is small enough to allow comparison between these different versions. The comparison can be also done by RTOS using the cost functions from Table 1 to make real time energy aware scheduling decisions.

The second class of parallel programs that our ECSA was applied to was streaming pipelines of communicating threads. There is a choice in how to spread the computation across threads to maximize throughput and therefore minimize execution time or lower the necessary device operating frequency. Having a number of available threads, a number of cores and the ability to apply voltage and frequency scaling, provides a wide range of configuration options in the design phase, with multiple optimization targets. This can range from optimizing for quality of service, time and energy, or a combination of all three. Our ECSA can take advantage of the fact that the energy model used can be parametric to voltage and frequency, to statically identify the most energy efficient configuration of the same program, among a number of different options that deliver the same required performance. The first step of analyzing the pipelined versions of industrial filter applications has been made in this paper. We are currently

working on extending our ECSA to automatically exploit the possible different configurations and provide the optimal solution, within the user's constraints.

Finally, the user needs to be aware of the potential effect of input data. When highly data sensitive applications are analyzed, the user can make some assumptions, based on the possible input data range, about the accuracy of the ECSA analysis. As explained in Figure 5, data that is close to random will lead to a smaller estimation error, when random data was used to build the energy model. From our findings, this variation can be up to 25%, but this has only been shown in short, contrived cases and is unlikely to be large in realistic programs.

5 Conclusion and Future Work

This work has given critical review of ECSA existing works that have overlooked the effect of using non data sensitive energy models and SRA bound techniques, on the retrieved energy estimations. In the absence of average case SRA and data sensitive energy models, we establish this effect in our experimental evaluation of ECSA on a set of mainly industrial benchmarks. We also demonstrate that such an analysis can still have a significant value for software developers, compiler engineers, development tools and RTOS, by establishing a number of ECSA applications in Section 4.3.

A technique was introduced to allow energy characterization of LLVM IR. It enables ECSA at this level with a small loss of accuracy, typically 1%, compared to ECSA at ISA level. ECSA is applied to a set of multi-threaded programs for the first time to our knowledge. This is a significant step beyond existing work that examines single-thread programs, because such an analysis can provide significant guidance for time-energy design space exploration between different numbers of threads and cores.

This work has generated new research questions. There is a clear need for non bounding SRA techniques that focus on average cases. Data sensitive energy models and SRA techniques are needed for ECSA to account for data sensitive switching activity in the processor. The majority of existing energy models are usually generated using random data. As we have discussed in Figure 5, alternative data energy models might be better for specific applications.

Future work aims to analyze more complex concurrent programs, such as distinct non-communicating threads rather than replicated threads, pipelines of threads with unbalanced workloads and other concurrency patterns. The ECSA can be combined with some more dynamic techniques such as abstract simulation to account for all the possible threads interleaving. Extending such analysis beyond deeply embedded systems, with more architectural performance enhancing features, might be done by exploiting more techniques from the WCET community, such as abstract interpretation and data cache analysis.

References

1. R. Jayaseelan, T. Mitra, and X. Li, "Estimating the worst-case energy consumption of embedded software," in *Real-Time and Embedded Technology and Applications*

Symposium, 2006. Proceedings of the 12th IEEE, pp. 81–90, April 2006.
2. V. Tiwari, S. Malik, A. Wolfe, and M. Tien-Chien Lee, "Instruction level power analysis and optimization of software," *Journal of VLSI Signal Processing Systems for Signal, Image, and Video Technology*, vol. 13, no. 2-3, pp. 223–238, 1996.
3. J. Pallister, S. Kerrison, J. Morse, and K. Eder, "Data dependent energy modelling: A worst case perspective," *CoRR*, vol. abs/1505.03374, 2015.
4. D. May, "XMOS XS1 Instruction Set Architecture," 2009.
5. S. Kerrison and K. Eder, "Energy Modeling of Software for a Hardware Multithreaded Embedded Microprocessor," *ACM Transactions on Embedded Computing Systems*, vol. 14, pp. 56:1–56:25, Apr. 2015.
6. C. Lattner and V. Adve, "LLVM: A compilation framework for lifelong program analysis and transformation," in *CGO*, pp. 75–88, 2004.
7. LLVMorg, "The LLVM Compiler Infrastructure," November 2014.
8. A. Bogliolo, L. Benini, G. Micheli, and B. Ricc, "Gate-Level Power and Current Simulation of CMOS Integrated Circuits," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 5, no. 4, pp. 473–488, 1997.
9. K. Roy and M. Johnson, "Software design for low power," in *Low power design in deep submicron electronics*, ch. 6, pp. 433–460, Kluwer Academic Publishers, 1997.
10. D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: A Framework for Architectural-Level Power Analysis and Optimizations," *ACM SIGARCH Computer Architecture News*, vol. 28, pp. 83–94, May 2000.
11. Sim-Panalyser, *Sim-Panalyser 2.0 Reference Manual*. 2004.
12. T. Austin, "SimpleScalar: An Infrastructure for computer system modeling," *IEEE Computer*, no. February, pp. 59–67, 2002.
13. S. Steinke, M. Knauer, L. Wehmeyer, and P. Marwedel, "An accurate and fine grain instruction-level energy model supporting software optimizations," in *Proc. of PATMOS*, Citeseer, 2001.
14. D. Sarta, D. Trifone, and G. Ascia, "A data dependent approach to instruction level power estimation," in *Low-Power Design, 1999. Proceedings. IEEE Alessandro Volta Memorial Workshop on*, pp. 182–190, Mar 1999.
15. H. Kojima, D. Gorny, K. Nitta, and K. Sasaki, "Power analysis of a programmable dsp for architecture/program optimization," in *Low Power Electronics, 1995., IEEE Symposium on*, pp. 26–27, Oct 1995.
16. M. Sami, D. Sciuto, C. Silvano, and V. Zaccaria, "An instruction-level energy model for embedded VLIW architectures," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 21, pp. 998–1010, Sept. 2002.
17. M. a. Ibrahim, M. Rupp, and H. Fahmy, "Power estimation methodology for VLIW Digital Signal Processors," in *2008 42nd Asilomar Conference on Signals, Systems and Computers*, no. 1, pp. 1840–1844, IEEE, Oct. 2008.
18. G. Qu, N. Kawabe, K. Usami, and M. Potkonjak, "Function-level power estimation methodology for microprocessors," *Proceedings of the 37th conference on Design automation - DAC '00*, pp. 810–813, 2000.
19. S. Lee, A. Ermedahl, and S. Min, "An Accurate Instruction-Level Energy Consumption Model for Embedded RISC Processors," *ACM SIGPLAN Notices*, vol. 36, pp. 1–10, Aug. 2001.
20. G. Contreras and M. Martonosi, "Power prediction for Intel XScale processors using performance monitoring unit events," in *ISLPED '05. Proceedings of the 2005 International Symposium on Low Power Electronics and Design*, pp. 221–226, IEEE, 2005.

21. Y. Shao and D. Brooks, "Energy characterization and instruction-level energy model of Intel's Xeon Phi processor," in *International Symposium on Low Power Electronics and Design (ISLPED)*, no. November, pp. 389–394, IEEE, Sept. 2013.
22. R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra, F. Mueller, I. Puaut, P. Puschner, J. Staschulat, and P. Stenström, "The worst-case execution-time problem—overview of methods and survey of tools," *ACM Trans. Embed. Comput. Syst.*, vol. 7, pp. 36:1–36:53, May 2008.
23. G. Brat, J. Navas, N. Shi, and A. Venet, "Ikos: A framework for static analysis based on abstract interpretation," in *Software Engineering and Formal Methods*, pp. 271–277, Springer, 2014.
24. B. Wegbreit, "Mechanical program analysis," *Commun. ACM*, vol. 18, no. 9, pp. 528–539, 1975.
25. M. Rosendahl, "Automatic complexity analysis," in *Proceedings of the Fourth International Conference on Functional Programming Languages and Computer Architecture*, FPCA '89, (New York, NY, USA), pp. 144–156, ACM, 1989.
26. S. K. Debray, N.-W. Lin, and M. Hermenegildo, "Task Granularity Analysis in Logic Programs," in *Proc. of the 1990 ACM Conf. on Programming Language Design and Implementation*, pp. 174–188, ACM Press, June 1990.
27. S. K. Debray, P. López-García, M. Hermenegildo, and N.-W. Lin, "Lower Bound Cost Estimation for Logic Programs," in *1997 International Logic Programming Symposium*, pp. 291–305, MIT Press, Cambridge, MA, October 1997.
28. P. Vasconcelos and K. Hammond, "Inferring Cost Equations for Recursive, Polymorphic and Higher-Order Functional Programs," in *Proceedings of the Workshop on Implementation of Functional Languages*, vol. 3145 of *Lecture Notes in Computer Science*, pp. 86–101, Springer-Verlag, September 2003.
29. J. Navas, E. Mera, P. López-García, and M. Hermenegildo, "User-Definable Resource Bounds Analysis for Logic Programs," in *International Conference on Logic Programming (ICLP'07)*, Lecture Notes in Computer Science, Springer, 2007.
30. E. Albert, P. Arenas, S. Genaim, and G. Puebla, "Closed-Form Upper Bounds in Static Cost Analysis," *Journal of Automated Reasoning*, vol. 46, pp. 161–203, February 2011.
31. J. Hoffmann, K. Aehlig, and M. Hofmann, "Multivariate amortized resource analysis," *ACM Trans. Program. Lang. Syst.*, vol. 34, no. 3, p. 14, 2012.
32. D. Alonso-Blas and S. Genaim, "On the limits of the classical approach to cost analysis," vol. 7460, pp. 405–421, 2012.
33. F. Zuleger, S. Gulwani, M. Sinn, and H. Veith, "Bound analysis of imperative programs with the size-change abstraction (extended version)," *CoRR*, vol. abs/1203.5303, 2012.
34. J. Engblom, A. Ermedahl, and F. Stappert, "Comparing different worst-case execution time analysis methods," in *The Work-in-Progress session of the 21st IEEE Real-Time Systems Symposium (RTSS 2000)*, November 2000.
35. Y.-T. Li and S. Malik, "Performance analysis of embedded software using implicit path enumeration," in *Proceedings of the 32Nd Annual ACM/IEEE Design Automation Conference*, DAC '95, (New York, NY, USA), pp. 456–461, ACM, 1995.
36. H. Theiling and C. Ferdinand, "Combining abstract interpretation and ilp for microarchitecture modelling and program path analysis," in *Real-Time Systems Symposium, 1998. Proceedings., The 19th IEEE*, pp. 144–153, Dec 1998.
37. G. Ottosson and M. Sjodin, "Worst-case execution time analysis for modern hardware architectures," in *In Proc. ACM SIGPLAN Workshop on Languages, Compilers and Tools for Real-Time Systems (LCT-RTS'97*, pp. 47–55, 1997.

38. D. Potop-Butucaru and I. Puaut, "Integrated Worst-Case Execution Time Estimation of Multicore Applications," in *13th International Workshop on Worst-Case Execution Time Analysis* (C. Maiza, ed.), vol. 30 of *OpenAccess Series in Informatics (OASIcs)*, (Dagstuhl, Germany), pp. 21–31, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2013.
39. P. Cousot and R. Cousot, "Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints," in *Conference Record of the Fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, (Los Angeles, California), pp. 238–252, ACM Press, New York, NY, 1977.
40. J. Navas, M. Méndez-Lojo, and M. Hermenegildo, "Safe Upper-bounds Inference of Energy Consumption for Java Bytecode Applications," in *The Sixth NASA Langley Formal Methods Workshop (LFM 08)*, April 2008. Extended Abstract.
41. M. Hermenegildo, G. Puebla, F. Bueno, and P. López-García, "Integrated Program Debugging, Verification, and Optimization Using Abstract Interpretation (and The Ciao System Preprocessor)," *Science of Computer Programming*, vol. 58, no. 1–2, 2005.
42. J. Navas, M. Méndez-Lojo, and M. Hermenegildo, "User-Definable Resource Usage Bounds Analysis for Java Bytecode," in *Proceedings of BYTECODE*, vol. 253 of *Electronic Notes in Theoretical Computer Science*, pp. 65–82, Elsevier - North Holland, March 2009.
43. U. Liqat, S. Kerrison, A. Serrano, K. Georgiou, P. Lopez-Garcia, N. Grech, M. Hermenegildo, and K. Eder, "Energy Consumption Analysis of Programs based on XMOS ISA-level Models," in *Proceedings of the 23rd International Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR'13)*, 2014.
44. D. Watt, *Programming XC on XMOS Devices*. XMOS Limited, 2009.
45. N. Grech, K. Georgiou, J. Pallister, S. Kerrison, J. Morse, and K. Eder, "Static analysis of energy consumption for llvm ir programs," in *Proceedings of the 18th International Workshop on Software and Compilers for Embedded Systems*, SCOPES '15, (New York, NY, USA), ACM, 2015.
46. P. Wägemann, T. Distler, T. Hönig, H. Janker, R. Kapitza, and W. Schröder-Preikschat, "Worst-case energy consumption analysis for energy-constrained embedded systems," July 2015.
47. K. Georgiou, "On the value and limits of multi-level energy consumption static analysis for deeply embedded single and multi-threaded programs - benchmarks." `https://www.cs.bris.ac.uk/home/kg8280/benchmarks.html`, 2015.
48. J. M. Townley, "Practical programming for static average-case analysis: the moqa investigation," 2013.
49. M. Schellekens, "Moqa; unlocking the potential of compositional static average-case analysis," *The Journal of Logic and Algebraic Programming*, vol. 79, no. 1, pp. 61 – 83, 2010. Speical Issue: Logic, Computability and Topology in Computer Science: A New Perspective for Old Disciplines.
50. F. Najm, "A survey of power estimation techniques in vlsi circuits," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 2, pp. 446–455, Dec 1994.
51. D. May, *The XMOS XS1 Architecture*. 2009.
52. K. Georgiou, "On the value and limits of multi-level energy consumption static analysis for deeply embedded single and multi-threaded programs - fnop modeling." `https://www.cs.bris.ac.uk/home/kg8280/fnops.html`, 2015.

53. C. Brandolese, S. Corbetta, and W. Fornaciari, "Software energy estimation based on statistical characterization of intermediate compilation code," in *Low Power Electronics and Design (ISLPED) 2011 International Symposium on*, pp. 333–338, Aug 2011.
54. "The dwarf debugging standard," Oct. 2013. `http://dwarfstd.org/`.
55. C. Lattner and D. Patel, "Extensible metadata in llvm ir," Apr 2014.
56. T. Wei, J. Mao, W. Zou, and Y. Chen, "A new algorithm for identifying loops in decompilation," in *Static Analysis* (H. Nielson and G. Fil, eds.), vol. 4634 of *Lecture Notes in Computer Science*, pp. 170–183, Springer Berlin Heidelberg, 2007.
57. Y.-S. Li and S. Malik, "Performance analysis of embedded software using implicit path enumeration," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 16, pp. 1477–1487, Dec 1997.
58. K. Eder, K. Georgiou, and N. Grech, eds., *Common Assertion Language*. ENTRA Project: Whole-Systems Energy Transparency (FET project 318337), November 2013. Deliverable 2.1, http://entraproject.eu.
59. XMOS, "xTimecomposer," November 2014.
60. M. Field, "Binary division," November 2014.
61. J. Hauser, "SoftFloat," November 2014.
62. XMOS, "Application Note: DSP performance on XS1-L device," November 2014.
63. J. Gustafsson, A. Betts, A. Ermedahl, and B. Lisper, "The mälardalen wcet benchmarks - past, present and future," in *Proceedings of the 10th International Workshop on Worst-Case Execution Time Analysis*, July 2010.
64. R. Wilhelm and D. Grund, "Computation takes time, but how much?," *Commun. ACM*, vol. 57, pp. 94–103, Feb. 2014.
65. C. Blackmore, O. Ray, and K. Eder, "A logic programming approach to predict effective compiler settings for embedded software," *Theory and Practice of Logic Programming*, vol. 15, pp. 481–494, 7 2015.
66. J. Pallister, K. Eder, and S. Hollis, "Optimizing the flash-ram energy trade-off in deeply embedded systems," in *Code Generation and Optimization (CGO), 2015 IEEE/ACM International Symposium on*, pp. 115–124, Feb 2015.
67. J. Pallister, K. Eder, S. J. Hollis, and J. Bennett, "A high-level model of embedded flash energy consumption," in *Proceedings of the 2014 International Conference on Compilers, Architecture and Synthesis for Embedded Systems*, CASES '14, (New York, NY, USA), pp. 20:1–20:9, ACM, 2014.
68. P. C. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *Proceedings of the 19th Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '99, (London, UK, UK), pp. 388–397, Springer-Verlag, 1999.

Acknowledgments

The research leading to these results has received funding from the European Union 7th Framework Programme (FP7/2007-2013) under grant agreement no 318337, ENTRA - Whole-Systems Energy Transparency. Special thanks to Intel for providing us with the equipment used for our power monitoring setup.